MIXTURE-OF-MICROZONES FOR SLOW AND FAST SENSORIMOTOR ADAPTATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Learning in naturalistic environments requires both rapid and long-term sensorimotor adaptation. How the brain and machines achieve such dual control is not clear. We introduce Mixture-of-Microzones (MoM), a cerebellar-inspired architecture that separates immediate error correction from slower learning. MoM uses error geometry to sparsely recruit specialized modules according to error direction and magnitude. Recruited modules perform local, state-conditioned coordinate transformations from sensory-error space to motor space, allowing their corrections to be composed and applied directly to the controller output. In parallel, the same gate determines which modules receive learning updates, progressively consolidating the response. In sensorimotor adaptation, this structured corrective pathway reduces adaptation time by approximately fourfold relative to ablated variants and substantially outperforms existing controller–adapter baselines. Ablations show that neither geometric gating nor a direct corrective pathway alone is sufficient: their combination is required to produce appropriately aligned motor corrections and rapid adaptation. We further isolate this principle in learned optimization, where directional gating generalizes to unseen coordinate transformations. Together, these results identify error-geometric coordinate transformation as a mechanism for correcting perturbations immediately while slower adaptive processes consolidate the response.

1 INTRODUCTION

Adaptive control requires solving two distinct problems: how to correct an error now, and how to learn from it for the future. Most learning-based approaches emphasize the latter, progressively modifying controller parameters. Yet unexpected perturbations often require correction before substantial learning can occur. This requires transforming an error observed in sensory or task space into an appropriate correction in action space:

$$\underbrace{\mathbf{e}_t}_{\text{sensory observed error}} \longrightarrow \underbrace{\Delta\mathbf{u}_t}_{\text{action-space correction}}$$

Biological motor control provides a striking example of this separation of timescales. The cerebellum supports both rapid error correction and slower sensorimotor adaptation, while cerebellar populations exhibit directional tuning to error (Herzfeld et al., 2015; 2018; Doostmohammadi et al., 2026). This suggests that error geometry may structure not only which components learn, but also how errors are transformed into immediate corrective actions Nguyen et al. (2026).

We introduce Mixture-of-Microzones (MoM), a cerebellar-inspired architecture built around this principle. Error direction and magnitude sparsely recruit specialized modules, each learning a local, state-conditioned transformation from sensory-error to motor coordinates. Their outputs are composed into an immediate correction to the nominal motor command, while the same geometric gate routes slower learning. In sensorimotor adaptation, MoM rapidly corrects unexpected perturbations by geometrically routing errors into appropriately aligned motor corrections. Ablations show that this benefit requires the combination of geometric gating and the direct corrective pathway. We further isolate this principle in learned optimization, where directional gating generalizes to unseen coordinate transformations.

Our contributions can be summarized as follows:

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1. We introduce a dual-timescale architecture that separates rapid error correction from slower adaptive learning.
2. We show that geometric gating structures local error-to-action coordinate transformations, enabling appropriately aligned corrections and rapid adaptation.
3. We demonstrate that this geometric routing principle generalizes to unseen rotations of gradient and actuator coordinates in learned optimization.

2 Related Work

Error geometry and cerebellar learning. Cerebellar learning signals are structured by the geometry of sensorimotor error. Purkinje cells exhibit preferred error directions, with climbing-fiber activity reflecting error direction and magnitude (Herzfeld et al., 2015; 2018; Shadmehr, 2020; Nguyen et al., 2026). Recent work further interprets this organization as a solution to credit assignment, where climbing-fiber signals reflect the projection of error onto the movement contribution, or *potent vector*, associated with individual Purkinje cells (Doostmohammadi et al., 2026). We build on this principle by using error geometry not only to assign learning, but also to route errors through local error-to-action transformations that generate immediate corrective actions.

Feedback-driven motor adaptation. Feedback can both correct ongoing behavior and drive slower learning. Recurrent controller–adapter models have shown how feedback signals can directly influence motor output while driving plasticity that progressively internalizes the correction (Filipe et al., 2026; Kaleb et al., 2024; Feulner et al., 2025; Pemberton et al., 2024). MoM extends this framework by structuring feedback through geometrically organized microzones: the same directional gate selects which local transformations contribute an immediate motor correction and which modules subsequently learn from the error.

3 Methods

3.1 Model Architecture

3.1.1 Controller–Adapter Architecture

We consider a closed-loop system comprising a continuous-time recurrent cortical controller and a cerebellar-inspired MoM adapter (Fig. 1). The cortical RNN models neural population dynamics with time constant τ, discretized via Euler integration step dt:

$$\mathbf{h}_{t+1} = \mathbf{h}_t + \frac{dt}{\tau}\left(-\mathbf{h}_t + W_h\Phi(\mathbf{h}_t) + W_{\text{in}}\mathbf{x}_t + W_{\text{fb}}\mathbf{cb}_t\right), \tag{1}$$

where $\Phi(\cdot)$ is a ReLU nonlinearity, $\mathbf{x}_t \in \mathbb{R}^3$ represents target-hold inputs, and $\mathbf{cb}_t$ is the modulatory cerebellar feedback projection. Network output corresponds to cartesian hand velocity $\mathbf{y}_t = W_{\text{out}}\Phi(\mathbf{h}_t) \in \mathbb{R}^2$. Desired position trajectory integrates online: $\tilde{\mathbf{p}}_{t+1} = \tilde{\mathbf{p}}_t + dt\,\mathbf{y}_t$. Inverse kinematics converts desired position into a nominal joint command $\mathbf{u}_t = \text{IK}(\tilde{\mathbf{p}}_t)$. Plant forward kinematics $\mathbf{p}_t = \text{FK}(\mathbf{u}_t^{\text{exec}})$ determine actual position, defining task-space error: $\mathbf{pe}_t = \mathbf{p}_t^* - \mathbf{p}_t$.

The MoM adapter contributes to control through two distinct pathways: (i) modulatory recurrent shaping via $\mathbf{cb}_t$, and (ii) direct joint-space error conversion via $\Delta\mathbf{u}_t$. Total executed joint commands combine additively:

$$\mathbf{u}_t^{\text{exec}} = \mathbf{u}_t + \Delta\mathbf{u}_t. \tag{2}$$

This architecture decouples slow, credit-assigned policy modulation in the recurrent loop from instant, state-coupled kinematic error correction.

3.1.2 Mixture-of-Microzones

The adapter consists of N expert modules, referred to as *microzones*. Each microzone i is associated with a fixed preferred direction $\boldsymbol{\phi}_i = \begin{bmatrix}\cos\alpha_i \\ \sin\alpha_i\end{bmatrix}$, $\qquad \alpha_i = \frac{2\pi i}{N}$,

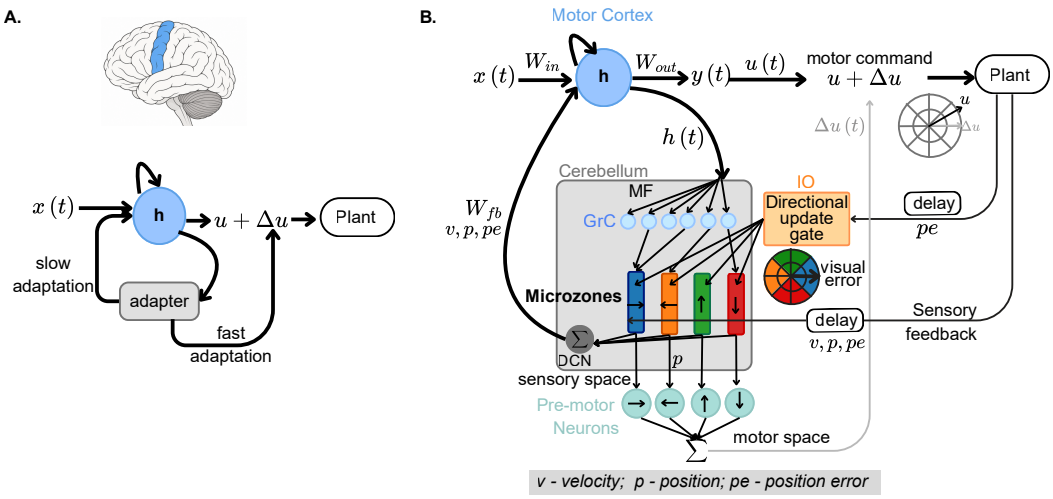

Figure 1: **Closed-loop cortico-cerebellar architecture for fast and slow motor adaptation. A.** Conceptual model in which the cortical controller h generates the motor command u, while a fast adaptive pathway provides an additive correction Δu. **B.** Neural implementation of the architecture. Cortical activity is expanded through cerebellar granule cells (GrCs) and distributed across microzones. Delayed position-error direction, conveyed through an inferior-olive (IO)-like teaching signal, selects the top-k microzones for adaptation. Microzone predictions follow two complementary pathways: in the recurrent pathway, predictions are integrated in a deep-cerebellar-nuclei (DCN)-like stage and fed back to the cortical recurrent neural network; in the direct pathway, microzone predictions undergo coordinated transformations to generate the adaptive correction Δu, which is added to the cortical output u to form the motor command $u + \Delta u$. DCN, deep cerebellar nuclei; IO, inferior olive; GrC, granule cells; MF, mossy fibers.

uniformly distributed over the two-dimensional error space. Rather than learning an unconstrained gating, MoM uses the geometry of the sensory error directly for expert recruitment.

For a delayed error vector $\mathbf{pe}_{t-\delta}$, the gating score of microzone i is

$$q_i(t) = \mathbf{pe}_{t-\delta}^{\top} \boldsymbol{\phi}_i = \underbrace{\|\mathbf{pe}_{t-\delta}\|}_{\text{error magnitude}} \cdot \underbrace{\cos(\theta_t - \alpha_i)}_{\text{directional alignment}} \tag{3}$$

where θ_t is the angular direction of $\mathbf{pe}_{t-\delta}$. This decomposition couples both error components: error direction determines which microzones are recruited, while error magnitude modulates recruitment sharpness.

We retain the k experts with the largest directional scores:

$$\mathcal{K}_t = \mathrm{TopK}(\mathbf{q}_t, k), \quad g_{\mathcal{K}_t} = \mathrm{softmax}(\mathbf{q}_{\mathcal{K}_t}), \quad g_i = 0 \quad \forall i \notin \mathcal{K}_t, \tag{4}$$

where $\mathbf{q}_t = [q_1(t), \ldots, q_N(t)]^{\top}$. For example, an error oriented at 45° is equally aligned with experts preferring 0° and 90°, resulting in equal routing weights ($g_1 = g_2 = 0.5$). Unless otherwise specified, we use $k = 2$, as two neighboring directional experts are sufficient to interpolate arbitrary directions in the 2D error space. Importantly, preferred routing directions are fixed rather than learned, imposing the geometry of the error space as an explicit inductive bias for modular specialization.

3.1.3 Cerebellar forward prediction path

The cerebellar adapter combines cortical activity with delayed sensorimotor feedback. The cortical state $\mathbf{h}_t$ is first projected through a $20\times$ mossy-fiber (MF) expansion Pemberton et al. (2024; 2021),

$$\mathbf{z}_t = \mathrm{ReLU}(W_{\mathrm{MF}} \mathbf{h}_t), \tag{5}$$

while delayed sensorimotor feedback is represented as $\mathbf{s}_{t-\delta} = [\mathbf{pe}_{t-\delta}; \mathbf{v}_{t-\delta}; \mathbf{p}_{t-\delta}]$, containing position error, velocity, and position.

Each microzone combines the expanded cortical representation with delayed feedback to predict a task outcome,

$$\mathbf{cb}_{i,t} = W_i \mathbf{z}_t + \boldsymbol{\gamma}_i \odot \mathbf{s}_{t-\delta}, \tag{6}$$

where W_i is the learned projection of microzone i, $\boldsymbol{\gamma}_i$ contains learned feedback gains, and $\odot$ denotes element-wise multiplication, such that each delayed sensory variable directly modulates only its corresponding prediction. The global cerebellar prediction is then

$$\mathbf{cb}_t = \sum_i g_i(t) \mathbf{cb}_{i,t}. \tag{7}$$

The adapter is trained to predict the desired task outcome $\mathbf{y}_t^* = [\mathbf{pe}_t^*; \mathbf{v}_t^*; \mathbf{p}_t^*]$, comprising the desired position error, velocity, and position:

$$\mathcal{L}_{\text{MoM}} = \|\mathbf{cb}_{t-\delta} - \mathbf{y}_t^*\|_2^2. \tag{8}$$

Because the mixture weights $g_i(t)$ also gate learning, error geometry determines both which microzones contribute to the prediction and which microzones are updated. The resulting cerebellar prediction is fed back to the recurrent controller through W_{fb}.

3.1.4 From error geometry to direct motor correction

Directional gating provides a natural mechanism for structuring rapid correction: error geometry determines which local error-to-action transformations are recruited, while error magnitude determines their contribution. To translate this structure into immediate behavior, we introduce a direct pathway from the cerebellar adapter to the motor command. This allows recruited microzones to correct the ongoing movement directly, rather than waiting for learning to update their predictive weights and indirectly influence the controller output.

However, the error and the motor command live in different coordinate spaces. The position error $\mathbf{pe}_t$ is expressed in cartesian sensory space, whereas the controller acts through the joint coordinates of the reaching arm. The direct pathway must therefore transform sensory corrections into motor corrections.

Rather than learning a single global sensory-to-motor mapping, MoM decomposes the error into local corrections, each assigned to the microzones most aligned with the error direction:

$$\Delta\mathbf{p}_{i,t} = g_i(t)\mathbf{pe}_t, \qquad i \in \mathcal{K}_t. \tag{9}$$

For example, a 45° error is shared between the neighboring 0° and 90° microzones. Each microzone therefore only needs to transform its local component of the correction into motor coordinates.

To perform this transformation generally across diverse effectors or abstract coordinate spaces without hardcoded kinematics, we replace analytical Jacobian inversion with a learned, state-conditioned local operator per microzone. Concatenating the localized error $\Delta\mathbf{p}_{i,t}$ with the microzone's predicted state $\hat{\mathbf{p}}_{i,t}$ from Sec. 3.1.3, the local motor-space correction is computed via a learnable transformation matrix $\mathbf{M}_i \in \mathbb{R}^{n_{\text{joints}} \times (D_{\text{error}} + D_{\text{state}})}$:

$$\Delta\mathbf{u}_{i,t} = \mathbf{M}_i \begin{bmatrix} \Delta\mathbf{p}_{i,t} \\ \hat{\mathbf{p}}_{i,t} \end{bmatrix}. \tag{10}$$

The local corrections are then summed and added to the nominal motor command,

$$\mathbf{u}_t^{\text{exec}} = \mathbf{u}_t + \sum_{i \in \mathcal{K}_t} \Delta\mathbf{u}_{i,t}. \tag{11}$$

This direct pathway enables rapid adaptation by composing posture-sensitive, localized sensorimotor transformations, while the predictive pathway provides the up-to-date state estimate required to condition these mappings despite delayed sensory feedback.

3.2 Experiments

3.2.1 Instructed-Delay Centre-Out Reaching Task

We construct a synthetic instructed-delay centre-out reaching task inspired by monkey experiments Perich et al. (2018). Initial and target positions $\mathbf{p}_{\text{start}}, \mathbf{p}_{\text{end}} \in \mathbb{R}^2$ are independently sampled

from $\mathcal{U}(-6, 6)$ per spatial axis. Desired velocity profiles are generated using a sigmoidal temporal profile:

$$f(t) = \frac{1}{1 + \exp(-\kappa t)}, \quad \kappa = 10\,\mathrm{s}^{-1}, \tag{12}$$

producing smooth, bell-shaped reaching kinematics scaled to target displacement.

Each trial spans 300 time steps ($\approx 3\,\mathrm{s}$) and includes a variable instructed-delay period sampled uniformly from $[1.2, 1.7]\,\mathrm{s}$. The hold constraint is encoded as a binary input channel via a Heaviside step indicator active throughout the delay window. Thus, the resulting task has a 3D input space $\mathbf{x}_t \in \mathbb{R}^3$ paired with 2D velocity targets $\mathbf{v}_t^* \in \mathbb{R}^2$.

Visuomotor Rotation (VR) Perturbation. To evaluate online error correction and adaptation, output projections are subject to a visuomotor rotation perturbation:

$$W_{\mathrm{out}}^{\mathrm{rot}} = R(\alpha) W_{\mathrm{out}}, \quad \text{where} \quad R(\alpha) = \begin{bmatrix} \cos\alpha & -\sin\alpha \\ \sin\alpha & \cos\alpha \end{bmatrix}, \tag{13}$$

introducing directional kinematics mismatch that requires online angular compensation by the network.

3.2.2 Training procedure

The recurrent controller comprises $N = 100$ units ($\tau = 50\,\mathrm{ms}$, $dt = 10\,\mathrm{ms}$). Cortical and cerebellar adaptation use distinct learning mechanisms. Following Feulner et al. (2025), cortical weights adapt through feedback-modulated eligibility traces, such that recent recurrent activity is reinforced according to the cerebellar feedback signal

$$\Delta W_{ji} \propto \sum_t \sum_k W_{\mathrm{fb},jk}\, cb_k(t)\, r_i(t), \tag{14}$$

where $r_i(t)$ is the eligibility trace of cortical unit i. In contrast, cerebellar learning is local to the microzones selected by the directional gate (Sec. 3.1.3), such that error geometry determines which experts are updated. Cerebellar parameters are optimized with Adam ($\eta = 10^{-3}$, $\beta = (0.9, 0.999)$).

Training proceeds in two phases. During **pre-training**, the model is first trained for 100 epochs with W_{fb} frozen, allowing the cerebellar adapter to develop its initial predictive representations. Pre-training then continues for 500 epochs with W_{fb} unfrozen, allowing cortical and cerebellar plasticity to operate jointly in closed loop.

During **adaptation**, cortical (RNN) weights and the transformation matrix are frozen and visuomotor perturbations $R(\alpha)$ are introduced. This isolates whether MoM can rapidly compensate for novel errors through its predictive and direct corrective pathways.

3.2.3 Generalization beyond sensorimotor control

Finally, we test whether the directional routing principle underlying MoM generalizes beyond the reaching task. Importantly, these experiments do not transfer the full MoM architecture; instead, we isolate its core inductive bias: fixed directional routing selects a sparse set of experts that learn local error-to-correction transformations.

We evaluate this principle in an additional domain using a *learned optimization* task, in which the optimizer must recover appropriate updates when observed gradients are rotated relative to the optimization coordinates. Models are trained on a small set of transformations and evaluated across held-out rotations, testing whether local directional experts can compose corrections outside their training conditions.

Full task definitions, model equations, training procedures, and baselines are provided in Appendix A.2.

4 RESULTS

4.1 GEOMETRY STRUCTURES RAPID ERROR CORRECTION

To test whether error geometry is necessary for effective correction, we independently ablated geometric gating and the direct pathway (Table 1). We quantify adaptation speed as the number of trials required for task error to return to its pre-perturbation baseline level (Appendix A.4). The full model adapted in 4 ± 1 trials, whereas removing either geometry or the direct pathway increased adaptation time to 16 ± 3 trials. Retaining the direct pathway without geometry provided almost no benefit over removing both components (16 ± 3 vs. 17 ± 4 trials).

To test whether geometry specifically enables the direct pathway to produce an appropriate motor correction, we measure its corrective alignment. For a direct joint-space correction $\Delta u^{\mathrm{direct}}$, we compute the cartesian displacement $\Delta p^{\mathrm{direct}} = f(u + \Delta u^{\mathrm{direct}}) - f(u)$, and measure its alignment with the task-space error e:

$$A = \cos(\Delta p^{\mathrm{direct}}, e).$$

Geometric gating increased corrective alignment from 0.24 ± 0.15 to 0.91 ± 0.05. Thus, a direct corrective pathway alone is insufficient: geometric gating structures the error-to-action transformation so that the resulting motor corrections are appropriately aligned with the task-space error.

Table 1: **Geometry enables effective direct correction.** Adaptation speed and task-space alignment of direct motor corrections (mean $\pm$ SD across 10 seeds).

Geometry	Direct path	Trials to adapt ↓	Corrective alignment ↑
✓	✓	$\mathbf{4_{\pm}2}$	$\mathbf{0.91_{\pm}0.05}$
✓	×	$16_{\pm}3$	–
×	✓	$16_{\pm}3$	$0.24_{\pm}0.15$
×	×	$17_{\pm}4$	–

4.2 MOM ENABLES RAPID ADAPTATION TO VISUOMOTOR PERTURBATIONS

We evaluate MoM under abrupt visuomotor rotations, using the cortico-cerebellar model of Pemberton et al. (2024) as baseline. MoM corrects trajectories within the first few trials after perturbation onset (Fig. 2A). Direct-pathway recruitment scales with perturbation magnitude (Fig. 2B), accelerating adaptation particularly for larger perturbations (Fig. 2C). Recruitment is transient, increasing sharply after perturbation onset and returning toward baseline as adaptation proceeds (Fig. 2D). Finally, learned local coordinate transformations match the performance of an analytical Jacobian following a 30° rotation, suppressing peak error within the first few trials (Fig. 2E). Importantly, comparable performance is achieved without requiring a closed-form kinematic model: the error-to-action transformation is learned directly rather than analytically specified. Appendix A.5 details the implementation of the analytical Jacobian baseline.

Finally, we extend the analysis beyond point-to-point reaching to complex MNIST drawing trajectories, where MoM preserves trajectory structure zero-shot under a 30° visuomotor rotation and further improves with adaptation (Appendix A.3).

4.3 ERROR GEOMETRY ORGANIZES MODULAR REPRESENTATIONS

We next examine how directional routing structures the learned experts. Because gating is determined by alignment with the error vector, nearby errors recruit overlapping experts, whereas opposing errors recruit distinct modules (Fig. 3A,B). With top-2 gating, intermediate directions are represented compositionally: for example, a 45° error is shared between neighboring 0° and 90° experts (Fig. 3C).

To test whether this organization arises specifically from the geometric routing rule, rather than from modularity alone, we compare MoM with an otherwise identical model in which the fixed directional gate is replaced by a learnable linear gate. This learned-gate model is free to discover its own routing structure from data, but produces substantially less directional specialization (Fig. 3D). In contrast, the structured organization induced by directional gating is preserved as the number of

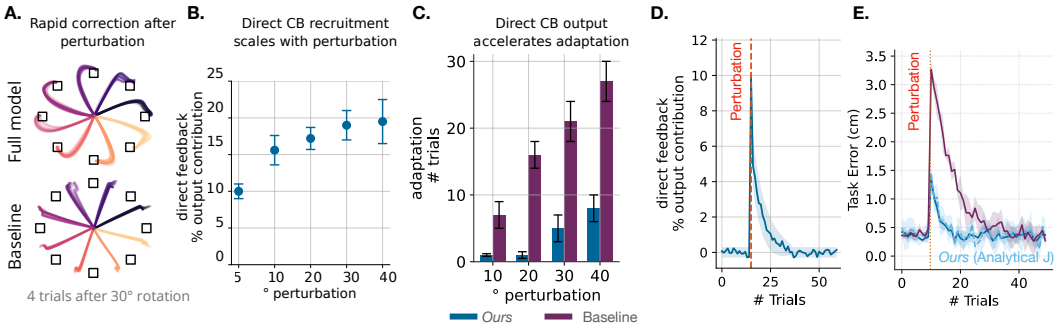

Figure 2: **Error-routed coordinate transformations enable rapid visuomotor correction.** **(A)** MoM rapidly corrects reaching trajectories following a 30° visuomotor rotation, compared with the cortico-cerebellar baseline Pemberton et al. (2024). **(B)** Direct-pathway recruitment increases with perturbation magnitude. **(C)** Direct correction reduces the number of trials required for adaptation, with larger benefits for larger perturbations. **(D)** Direct-pathway recruitment peaks following perturbation onset and declines as adaptation proceeds. **(E)** Learned local mappings match analytical-Jacobian performance without requiring a closed-form kinematic model. Shaded regions denote ± 1 SD ($n = 10$ seeds).

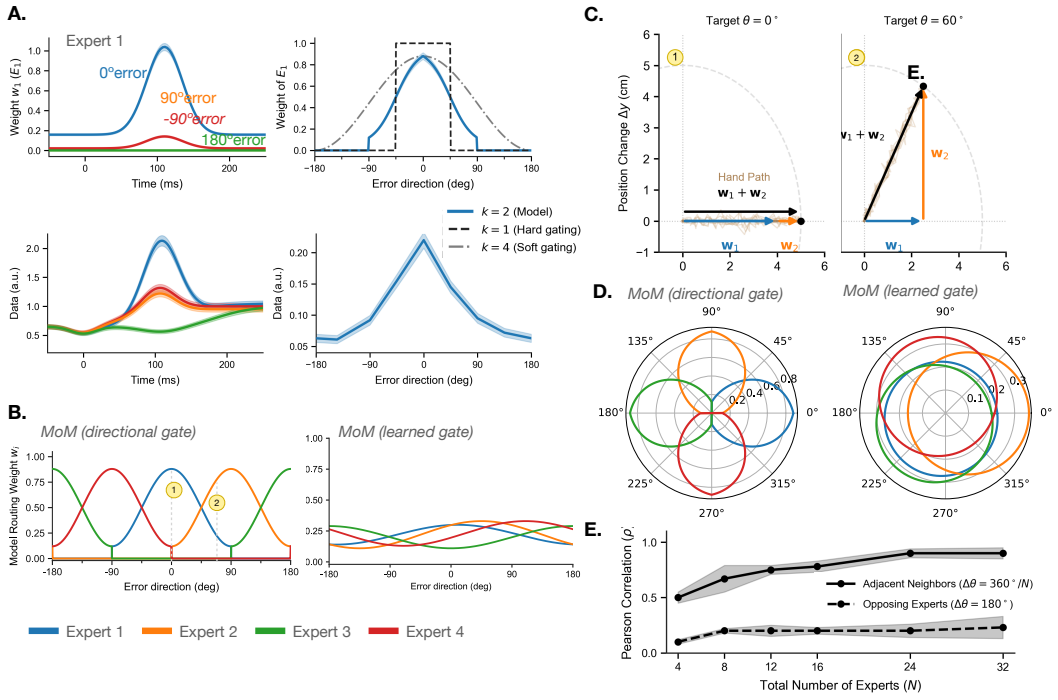

Figure 3: **Error geometry organizes modular representations.** **(A)** Directional tuning of a single model expert (*top*) and an experimentally recorded neuron (*bottom*; adapted from Herzfeld et al. (2018)). Responses are maximal for errors along the preferred direction, weaker for neighboring ($\pm 90°$) directions, and minimal for the opposing (180°) direction. **(B)** Nearby error directions recruit overlapping experts, whereas opposing directions recruit distinct modules. **(C)** Top-2 gating compositionally represents intermediate directions through neighboring experts. **(D)** Geometric gating produces structured expert-specific directional tuning, which is largely absent with a learned gate. **(E)** Expert organization across different numbers of directional microzones. Bands indicate ± 1 SD across $n = 10$ runs.

experts is varied (Fig. 3E). Importantly, this structure is behaviorally consequential: top-2 geometric routing enables substantially faster correction than hard, fully soft, or learned gating, while also

Table 2: **Geometric routing balances expert specialization and compositionality for rapid adaptation.** Adjacent- and opposing-expert correlations (ρ_{adj}, ρ_{opp}), trials to behavioral correction (T_{corr}), trials to internalization (T_{int}), and zero-shot task error. Metrics definitions are provided in Appendix A.4. Values are mean ± SD across 10 seeds.

Model Variant	$\rho_{\text{adj}} \uparrow$	$\rho_{\text{opp}} \downarrow$	$\mathbf{T}_{\text{corr}} \downarrow$	$\mathbf{T}_{\text{int}} \downarrow$	**Zero-shot Error (cm)** $\downarrow$
Filipe et al. (2026)	–	–	$16_{\pm 3}$	$24_{\pm 3}$	$2.1_{\pm 0.3}$
Pemberton et al. (2024)	–	–	$21_{\pm 4}$	$21_{\pm 3}$	$3.4_{\pm 0.4}$
Ours ($K = 2, N = 4$)	$\mathbf{0.50}_{\pm \mathbf{0.05}}$	$\mathbf{0.11}_{\pm \mathbf{0.02}}$	$\mathbf{4}_{\pm \mathbf{2}}$	$25_{\pm 4}$	$\mathbf{1.5}_{\pm \mathbf{0.3}}$
Gate strategy					
Hard Gate ($K = 1$)	$0.38_{\pm 0.04}$	$0.20_{\pm 0.03}$	$14_{\pm 3}$	$25_{\pm 3}$	$1.8_{\pm 0.2}$
Soft Gate ($K = 4$)	$0.65_{\pm 0.05}$	$0.40_{\pm 0.03}$	$16_{\pm 3}$	$25_{\pm 2}$	$1.5_{\pm 0.4}$
Learned Gate	$0.80_{\pm 0.08}$	$0.75_{\pm 0.12}$	$21_{\pm 3}$	$25_{\pm 4}$	$3.0_{\pm 0.4}$

improving zero-shot performance relative to the learned gate (Table 2). Together, these results show that error geometry organizes experts into a structured, compositional representation that supports rapid correction and adaptation.

4.4 Direct correction and learned feedback support distinct adaptation timescales

We next examine how MoM's pathways contribute to behavioral and neural adaptation, comparing model responses with experimental data (black), Perich et al. (2018). Direct correction primarily reduces task error early after perturbation, with differences diminishing later in adaptation (Fig. 4A). Interestingly, the baseline adapts its take-off error faster than both MoM and the experimental data (Fig. 4B). Without direct correction, larger residual task errors place greater pressure on the internal adapter to adapt, whereas MoM can correct behavior directly while its internal representation adapts more gradually.

Population activity further reveals two within-trial components (Fig. 4C): an initial forward-prediction peak followed by a delayed sensory-feedback-dependent peak. The latter is absent in the baseline and attenuated when directional gating or sensory feedback is ablated. Together, these results support distinct contributions of rapid feedback-based correction and slower internal adaptation.

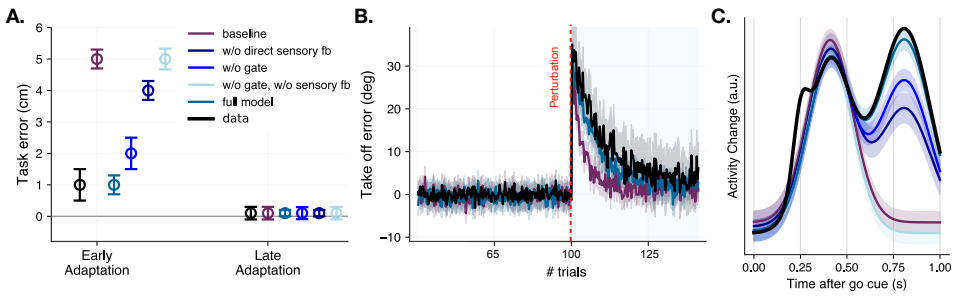

Figure 4: **Direct correction and learned feedback support distinct adaptation timescales.** Model variants are compared with experimental data (black) Perich et al. (2018). **(A)** Task error during early and late adaptation. **(B)** Take-off error across adaptation; MoM follows the slower time course observed experimentally. **(C)** Within-trial population activity reveals an initial forward-prediction peak and a delayed feedback-dependent peak, which is attenuated by feedback-related ablations. Lines and points indicate the mean; shaded regions and error bars represent the standard error of the mean across 10 random seeds.

4.5 Error-geometric routing generalizes beyond sensorimotor control

We next isolate geometric gating from the full sensorimotor architecture. We construct a learned optimization problem in which the optimizer observes a rotated rather than the true gradient, and must

learn to transform it into an appropriate update. Four directional experts learn local gradient-to-update transformations using only three training rotations ($-90°$, $0°$, and $90°$).

Despite this sparse training, directional gating generalizes across training horizons and objective geometries, including a rotated quadratic that breaks alignment with the expert directions (Fig. 5A). The advantage persists across unseen rotations (Fig. 5B), where learned and random gating degrade away from training conditions. Under a held-out $45°$ rotation, neighboring experts compose their local transformations to recover efficient optimization trajectories (Fig. 5C). Thus, geometric gating supports generalization of learned error-to-update mappings beyond the transformations seen during training.

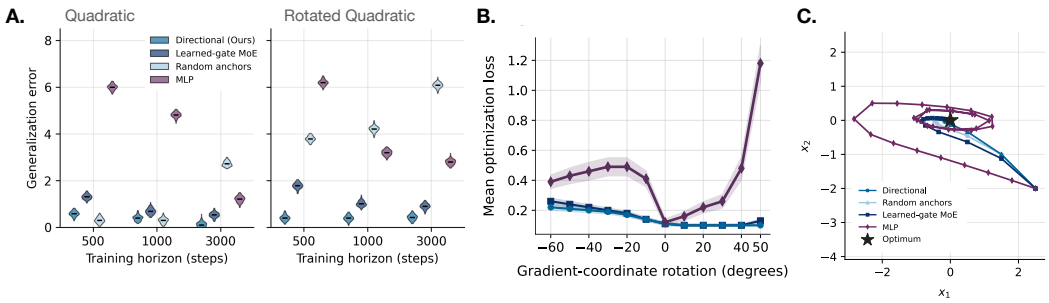

Figure 5: **Geometric routing generalizes to unseen rotations in learned optimization. (A)** Generalization across training horizons and objective geometries. **(B)** Optimization loss under unseen gradient-coordinate rotations. **(C)** Optimization trajectories under a held-out $45°$ rotation. Lines and points indicate the mean; shaded regions and error bars represent the standard error of the mean across 15 random seeds.

5 Discussion

Beyond faster adaptation, MoM suggests a functional role for the geometry of error representations. In neuroscience, directional cerebellar error tuning is typically associated with learning and credit assignment. Our results suggest an additional role: the same organization can determine which local error-to-action transformations contribute to an immediate correction. This provides a computational link between structured cerebellar error signals, rapid feedback-dependent correction, and slower internalization in the adapter dynamics.

For machine learning and control, this principle may be useful whenever the mapping between observed errors and corrective actions changes across conditions. Examples include robotic manipulation under changing tool or payload dynamics, adaptive control under actuator perturbations, and learned optimizers operating across changing parameterizations. Rather than relearning a global controller for each condition, geometric routing can recruit and compose local error-to-action transformations, providing a structured inductive bias for rapid adaptation and generalization.

Limitations and future work. Our experiments are controlled and low-dimensional, serving as a proof of concept for geometric error routing. The current model assumes that a meaningful error geometry is available a priori; extending this principle to higher-dimensional and less explicitly structured error spaces remains an important challenge. Future work should also test MoM on more complex tasks and investigate whether the adapter can complement richer policies and control architectures beyond the RNN considered here. Such settings will determine whether geometric error routing can provide rapid correction while preserving the capabilities of larger learned controllers.

AI use statement

In this work, we used generative AI tools for language editing and rephrasing of author-written text. We have not used generative AI tools for generating research ideas, developing the methodology, writing experimental code, producing experimental results, or generating figures, and other required

disclosure tasks are not applicable to this work. Additionally, we used generative AI tools for improving the clarity, conciseness, grammar, and presentation of the text. We have reviewed all AI-assisted work. All AI-assisted revisions were manually reviewed by the authors to ensure that they accurately reflected the intended scientific meaning and did not alter the methods, experimental results, or scientific claims. All research ideas, methodological decisions, experiments, analyses, and figures were developed and produced by the authors. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

Reproducibility statement

We provide detailed descriptions of the model architecture, training procedure, and experimental setup in the main text. Appendix A provides full definitions of the optimization benchmark, together with evaluation metrics and implementation details, while Appendix B describes the analytical Jacobian baseline. Hyperparameters and training conditions are reported throughout the Methods and Appendix. All results are reported across multiple random seeds, with variability indicated in the corresponding figures and tables.

References

Luca Agueci and N. A. Cayco-Gajic. Distributed learning across fast and slow neural systems supports efficient motor adaptation. *bioRxiv*, 2025. doi: 10.1101/2025.06.01.657238.

James S. Albus. A theory of cerebellar function. *Mathematical Biosciences*, 10(1-2):25–61, 1971. doi: 10.1016/0025-5564(71)90051-4.

Amy J. Bastian. Learning to predict the future: the cerebellum adapts feedforward movement control. *Current Opinion in Neurobiology*, 16(6):645–649, 2006. doi: 10.1016/j.conb.2006.08.016.

François G. C. Blot, Joshua J. White, Amy van Hattem, Licia Scotti, Vaishnavi Balaji, Youri Adolfs, R. Jeroen Pasterkamp, Chris I. De Zeeuw, and Martijn Schonewille. Purkinje cell microzones mediate distinct kinematics of a single movement. *Nature Communications*, 14(1):4358, 2023. doi: 10.1038/s41467-023-40111-5. URL https://doi.org/10.1038/s41467-023-40111-5.

Lior Botzer and Amir Karniel. Feedback and feedforward adaptation to visuomotor delay during reaching and slicing movements. *European Journal of Neuroscience*, 38(5):2431–2441, 2013. doi: 10.1111/ejn.12211.

Ellen J. P. Boven, Joseph O. Pemberton, Paul Chadderton, Richard Apps, and Rui Ponte Costa. Cerebro-cerebellar networks facilitate learning through feedback decoupling. *Nature Communications*, 14(1):51, 2023. doi: 10.1038/s41467-022-35658-8. URL https://doi.org/10.1038/s41467-022-35658-8.

Chris I. De Zeeuw. Bidirectional learning in upbound and downbound microzones of the cerebellum. *Nature Reviews Neuroscience*, 22:92–110, 2021. doi: 10.1038/s41583-020-00392-x. URL https://doi.org/10.1038/s41583-020-00392-x.

Jafar Doostmohammadi, Nazanin Mohammadrezaei, Hisham Y. Elseweifi, Alana D. Chandler, Elijah A. Taeckens, and Reza Shadmehr. Climbing fibers encode the gradient of a loss function for the cerebellum. *bioRxiv*, 2026. doi: 10.64898/2026.07.27.741034. URL https://doi.org/10.64898/2026.07.27.741034.

Mohammad Amin Fakharian, Alden M. Shoup, Paul Hage, Hisham Y. Elseweifi, and Reza Shadmehr. A vector calculus for neural computation in the cerebellum. *Science*, 388(6749):869–875, 2025. doi: 10.1126/science.adu6331. URL https://doi.org/10.1126/science.adu6331.

Barbara Feulner, Matthew G. Perich, Lee E. Miller, Claudia Clopath, and Juan A. Gallego. A neural implementation model of feedback-based motor learning. *Nature Communications*, 16 (1):1805, 2025. doi: 10.1038/s41467-024-54738-5. URL https://doi.org/10.1038/s41467-024-54738-5.

Ana Carolina Filipe, Rui Ponte Costa, and Cláudia Soares. Error- and prediction-driven motor learning in the cortico-cerebellar loop, 2026. URL https://arxiv.org/abs/2609.29945.

Kim M. Gruver, Jenny W. Y. Jiao, Eviatar Fields, Sen Song, Per Jesper Sjöström, and Alanna J. Watt. Structured connectivity in the output of the cerebellar cortex. *Nature Communications*, 15(1):5563, 2024. doi: 10.1038/s41467-024-49339-1. URL https://doi.org/10.1038/s41467-024-49339-1.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015.

David J. Herzfeld, Yoshiko Kojima, Robijanto Soetedjo, and Reza Shadmehr. Encoding of action by the Purkinje cells of the cerebellum. *Nature*, 526:439–442, 2015. doi: 10.1038/nature15693. URL https://doi.org/10.1038/nature15693.

David J. Herzfeld, Yoshiko Kojima, Robijanto Soetedjo, and Reza Shadmehr. Encoding of error and learning to correct that error by the purkinje cells of the cerebellum. *Nature Neuroscience*, 21:736–743, 2018. doi: 10.1038/s41593-018-0136-y. URL https://doi.org/10.1038/s41593-018-0136-y.

Masao Ito. Bases and implications of learning in the cerebellum—adaptive control and internal model mechanism. *Progress in Brain Research*, 148:95–109, 2005. doi: 10.1016/S0079-6123(04)48008-7.

Jun Izawa and Reza Shadmehr. Learning from sensory and reward prediction errors during motor adaptation. *PLoS Computational Biology*, 7(3):e1002012, 2011. doi: 10.1371/journal.pcbi.1002012.

Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. Adaptive mixtures of local experts. *Neural Computation*, 3(1):79–87, 1991. doi: 10.1162/neco.1991.3.1.79.

K. Kaleb, B. Feulner, J. A. Gallego, and C. Clopath. Feedback control guides credit assignment in recurrent neural networks. In *Advances in Neural Information Processing Systems*, 2024.

Mitsuo Kawato. Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 5(6):718–727, 1995. doi: 10.1016/0959-4388(95)80044-5.

Maedbh King, Carlos R. Hernandez-Castillo, Russell A. Poldrack, Richard B. Ivry, and Jörn Diedrichsen. Functional boundaries in the human cerebellum revealed by a multi-domain task battery. *Nature Neuroscience*, 22(8):1371–1378, 2019. doi: 10.1038/s41593-019-0436-x. URL https://doi.org/10.1038/s41593-019-0436-x.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

J. W. Krakauer, Z. M. Pine, M.-F. Ghilardi, and C. Ghez. Learning of visuomotor transformations for vectorial planning of reaching trajectories. *Journal of Neuroscience*, 2000.

David Marr. A theory of cerebellar cortex. *The Journal of Physiology*, 202(2):437–470, 1969.

V. Nguyen, C. Gros, and B. M. Stell. Rapid motor skill adjustment is associated with population-level modulation of cerebellar error signals. *Nature Neuroscience*, 2026. doi: 10.1038/s41593-025-02126-7. URL https://doi.org/10.1038/s41593-025-02126-7.

Marco Novello, Laurens W. J. Bosman, and Chris I. De Zeeuw. A systematic review of direct outputs from the cerebellum to the brainstem and diencephalon in mammals. *The Cerebellum*, 23(1):210–239, 2024. doi: 10.1007/s12311-022-01499-w. URL https://doi.org/10.1007/s12311-022-01499-w.

Joseph O. Pemberton, Ellen J. P. Boven, Richard Apps, and Rui Ponte Costa. Cortico-cerebellar networks as decoupling neural interfaces. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

Joseph O. Pemberton, Paul Chadderton, and Rui Ponte Costa. Cerebellar-driven cortical dynamics can enable task acquisition, switching and consolidation. *Nature Communications*, 15:10913, 2024. doi: 10.1038/s41467-024-55315-6. URL https://doi.org/10.1038/s41467-024-55315-6.

Matthew G. Perich, Juan A. Gallego, and Lee E. Miller. A neural population mechanism for rapid learning. *Neuron*, 2018.

Liviu S. Popa, Adam L. Hewitt, and Timothy J. Ebner. The cerebellum for jocks and nerds alike. *Frontiers in Systems Neuroscience*, 13:71, 2019. doi: 10.3389/fnsys.2019.00071.

Jennifer L. Raymond and Javier F. Medina. Computational principles of supervised learning in the cerebellum. *Annual Review of Neuroscience*, 41:233–253, 2018. doi: 10.1146/annurev-neuro-080317-061948.

Reza Shadmehr. Population coding in the cerebellum: a machine learning perspective. *Journal of Neurophysiology*, 2020. doi: 10.1152/jn.00449.2020. URL https://doi.org/10.1152/jn.00449.2020.

Reza Shadmehr, Maurice A. Smith, and John W. Krakauer. Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience*, 33:89–108, 2010. doi: 10.1146/annurev-neuro-060909-153135.

Ludovic Spaeth, Jyotika Bahuguna, Theo Gagneux, Kevin Dorgans, Izumi Sugihara, Bernard Poulain, Demian Battaglia, and Philippe Isope. Cerebellar connectivity maps embody individual adaptive behavior in mice. *Nature Communications*, 13:580, 2022. doi: 10.1038/s41467-022-27984-8. URL https://doi.org/10.1038/s41467-022-27984-8.

David Sussillo, Mark M. Churchland, Matthew T. Kaufman, and Krishna V. Shenoy. A neural network that finds a naturalistic solution for the production of muscle activity. *Nature Neuroscience*, 18(7):1025–1033, 2015. doi: 10.1038/nn.4042.

Jordan A. Taylor and Richard B. Ivry. The role of strategies in motor learning. *Annals of the New York Academy of Sciences*, 2012. doi: 10.1111/j.1749-6632.2011.06430.x.

Antoine M. Valera, Francesca Binda, Sophie A. Pawlowski, Jean-Luc Dupont, Jean-François Casella, Jeffrey D. Rothstein, Bernard Poulain, and Philippe Isope. Stereotyped spatial patterns of functional synaptic connectivity in the cerebellar cortex. *eLife*, 2016. doi: 10.7554/eLife.09862. URL https://doi.org/10.7554/eLife.09862.

A Appendix

A.1 Implementation details

Table 3 summarizes the main implementation and training parameters for the sensorimotor experiments.

Plant implementation. The reaching plant is implemented as a planar two-link arm. The cortical RNN outputs a cartesian hand velocity,

$$\mathbf{y}_t = W_{\text{out}}\Phi(\mathbf{h}_t),$$

which is integrated to obtain the desired cartesian hand position,

$$\tilde{\mathbf{p}}_{t+1} = \tilde{\mathbf{p}}_t + dt\,\mathbf{y}_t.$$

Because the RNN output is expressed in cartesian task space whereas the plant is controlled in joint space, inverse kinematics is applied at every time step to obtain the nominal joint command,

$$\mathbf{u}_t = \text{IK}(\tilde{\mathbf{p}}_t).$$

The MoM direct pathway produces an additional joint-space correction $\Delta\mathbf{u}_t$, giving the executed command

$$\mathbf{u}_t^{\text{exec}} = \mathbf{u}_t + \Delta\mathbf{u}_t.$$

The plant state is then obtained using the two-link arm forward kinematics,

$$\mathbf{p}_t = \text{FK}(\mathbf{u}_t^{\text{exec}}),$$

Table 3: **Main model and training hyperparameters.**

Parameter	Value
Cortical controller	
Number of recurrent units	100
Activation	ReLU
Time constant τ	50 ms
Integration step dt	10 ms
Input dimension	3
Output dimension	2
Trial duration	300 steps ($\approx$ 3 s)
Instructed delay	$\mathcal{U}(1.2, 1.7)$ s
Mixture-of-Microzones	
Active microzones K	2
Preferred directions	Uniformly spaced in $[0, 2\pi)$
Sensorimotor delay δ	100 ms
Mossy-fiber expansion factor	$20\times$
Expanded cortical dimension z_t	2000
Number of microzones N	4
Training	
Cerebellar optimizer	Adam
Cerebellar learning rate	10^{-3}
Adam (β_1, β_2)	$(0.9, 0.999)$
Pre-training, frozen W_{fb}	100 epochs
Pre-training, joint plasticity	500 epochs
Cortical weights during adaptation	Frozen
Number of random seeds	10
Perturbation / evaluation	
Visuomotor perturbation	Rotation of W_{out}
Perturbation onset	After convergence to near-zero baseline error (trial 150 in reported experiments)
Default perturbation angle	30°
Plant	
Plant	Planar two-link arm
Link length l_1	6
Link length l_2	10
Plant input	Joint configuration $\mathbf{u}_t \in \mathbb{R}^2$
Plant output	Cartesian position $\mathbf{p}_t \in \mathbb{R}^2$

and the resulting Cartesian position is used to compute the sensory task error $\mathbf{e}_t = \mathbf{p}_t^\star - \mathbf{p}_t$. Thus, inverse kinematics converts the nominal cartesian trajectory generated by the RNN into joint-space commands, while the MoM pathway learns the local sensory-to-joint coordinate transformation required for direct correction.

For joint configuration $\boldsymbol{\theta} = [\theta_1, \theta_2]^\top$, the plant forward kinematics are

$$\mathrm{FK}(\boldsymbol{\theta}) = \begin{bmatrix} l_1 \cos\theta_1 + l_2 \cos(\theta_1 + \theta_2) \\ l_1 \sin\theta_1 + l_2 \sin(\theta_1 + \theta_2) \end{bmatrix}$$

with link lengths l_1 and l_2 (Table 3).

A.2 Generalization benchmarks

To test whether error-geometric gating captures a general computational principle rather than a mechanism specific to reaching, we transfer only the directional gating structure of MoM to two additional domains. These experiments do not use the recurrent controller, predictive cerebellar pathway, or direct kinematic correction of the full model.

A.2.1 Optimization under rotated gradients

We consider a two-dimensional anisotropic quadratic objective,

$$f(\mathbf{x}) = \frac{1}{2}\mathbf{x}^\top H\mathbf{x}, \qquad H = \mathrm{diag}(1, 10), \tag{15}$$

where $\mathbf{x} \in \mathbb{R}^2$ is the optimization state and H determines the curvature of the objective. We additionally consider a 30°-rotated objective, $H_{\mathrm{rot}} = QHQ^\top$, where Q is the corresponding rotation matrix, to test whether routing depends on alignment with the coordinate axes.

At optimization step t, the true gradient is $\mathbf{g}_t = \nabla f(\mathbf{x}_t)$. Rather than observing $\mathbf{g}_t$ directly, the optimizer receives a coordinate-rotated gradient, $\tilde{\mathbf{g}}_t = R_\theta \mathbf{g}_t$, where $R_\theta \in \mathbb{R}^{2\times 2}$ rotates the gradient by perturbation angle θ. Thus, $\tilde{\mathbf{g}}_t$ preserves gradient magnitude but changes the mapping between the observed descent direction and the update required in the original coordinate system.

We instantiate four uniformly spaced directional experts with Top-2 gating. For each observed descent direction $-\tilde{\mathbf{g}}_t$, the gate selects the expert set $\mathcal{K}_t$ and assigns routing weights $g_i(t)$. Each expert learns a local linear transformation $M_i \in \mathbb{R}^{2\times 2}$, and their outputs are composed as

$$\Delta\mathbf{x}_t = \eta \sum_{i\in\mathcal{K}_t} g_i(t) M_i(-\tilde{\mathbf{g}}_t), \qquad \mathbf{x}_{t+1} = \mathbf{x}_t + \Delta\mathbf{x}_t, \tag{16}$$

where $\Delta\mathbf{x}_t$ is the learned optimization update and η is a learned global step size.

Models are meta-trained by unrolling 20 optimization steps using only $\theta \in \{-90^\circ, 0^\circ, 90^\circ\}$ and evaluated on unseen rotations. We compare directional routing against random directional anchors, a learned-gate MoE, and a dense MLP optimizer.

A.3 Advanced Sensorimotor Task

To evaluate MoM on more complex motor trajectories, we adopt the MNIST-based drawing task introduced by (Pemberton et al., 2024). In this task, the model receives an MNIST digit as input and is trained to generate a corresponding two-dimensional drawing trajectory. We extend this setting by introducing an abrupt 30° visuomotor rotation after the model has learned the unperturbed task. This perturbation tests whether MoM can adapt previously learned, temporally structured movements containing continuously varying directions and curvature. We evaluate trajectories before perturbation, immediately after perturbation (zero-shot), and following adaptation to the 30° rotation (Fig. 6).

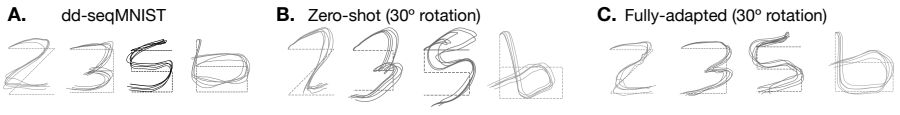

Figure 6: **MoM adapts complex drawing trajectories under visuomotor perturbations. A.** Model-generated trajectories for representative dd-seqMNIST digits under the unperturbed condition; dashed lines indicate target trajectories and solid lines indicate model outputs. **B.** Zero-shot trajectories immediately following the introduction of a 30° visuomotor rotation. **C.** Trajectories after adaptation to the 30° rotation, demonstrating recovery of the target digit trajectories despite their continuously varying movement directions and curvature.

MoM preserves the overall structure of the digit trajectories even zero-shot, immediately following the unseen 30° rotation (Fig. 6B), despite substantial changes in the required error-to-action mapping. Following adaptation, the trajectories recover toward their target shapes (Fig. 6C). These results indicate that error-geometric routing extends beyond simple point-to-point reaches, supporting both immediate correction and subsequent adaptation of more complex motor trajectories.

A.4 Evaluation Metrics

Trials to behavioral correction (T_{corr}) We measure adaptation speed as the number of trials after perturbation onset required for task error to return to its pre-perturbation baseline level. The baseline

error E_{base} is computed as the mean task error over the final 10 trials preceding perturbation. A model is considered adapted at the first trial t for which task error remains at or below E_{base} for 10 consecutive trials:

$$T_{\text{corr}} = \min \{t : E_{t:t+K-1} \leq E_{\text{base}}\}.$$

Corrective alignment. We quantify whether the direct motor correction produces movement in the appropriate task-space direction. Let u_t denote the nominal joint-space motor command and $\Delta u_t^{\text{direct}}$ the correction generated by the direct pathway. Using the arm forward kinematics $f(\cdot)$, the cartesian displacement induced by this correction is

$$\Delta p_t^{\text{direct}} = f(u_t + \Delta u_t^{\text{direct}}) - f(u_t).$$

Corrective alignment is the cosine similarity between this displacement and the cartesian task error $e_t = p_t^\star - p_t$:

$$A_t = \frac{(\Delta p_t^{\text{direct}})^\top e_t}{\|\Delta p_t^{\text{direct}}\|_2 \|e_t\|_2}.$$

We compute A_t over the final 100 ms of each 300 ms trial and average over this window. An alignment of 1 indicates correction toward the target, 0 an orthogonal correction, and -1 correction away from the target. We report mean $\pm$ SD across 10 random seeds.

Trials to internalization (T_{int}). We quantify internal adaptation using take-off error. Take-off error is the angular difference between the straight-line direction from the starting position to the target and the initial movement direction, estimated from the first 10 trajectory points after movement onset. This measure captures the initial motor plan before later feedback-dependent corrections substantially alter the trajectory. T_{int} is the number of trials after perturbation onset required for take-off error to reach the straight-to-target direction.

Adjacent- and opposing-expert correlation (ρ_{adj}, ρ_{opp}). We quantify the geometric organization of the expert representations using pairwise correlations between expert responses across error directions. ρ_{adj} is the mean correlation between experts with neighboring preferred directions, whereas ρ_{opp} is the mean correlation between experts with opposing preferred directions. High ρ_{adj} indicates shared representation between neighboring experts, while low ρ_{opp} indicates separation between experts representing opposing error directions.

Zero-shot error. Zero-shot error measures task performance immediately after the introduction of a perturbation, before any adaptation to that perturbation has occurred. It is defined as the task error on the first trial following perturbation onset. Lower values therefore indicate better immediate generalization to the new perturbation.

A.5 Closed-form kinematics Oracle

To establish a physics-engineered baseline for the 2D reaching tasks, we implement an analytical Jacobian inversion model. For a planar arm with link lengths l_1, l_2 and joint angles $\boldsymbol{\theta} = [\theta_1, \theta_2]^\top$, the forward kinematics mapping end-effector position $\mathbf{p} = [x, y]^\top$ is given by:

$$\mathbf{p}(\boldsymbol{\theta}) = \begin{bmatrix} l_1 \cos\theta_1 + l_2 \cos(\theta_1 + \theta_2) \\ l_1 \sin\theta_1 + l_2 \sin(\theta_1 + \theta_2) \end{bmatrix}. \tag{17}$$

Differentiation yields the 2×2 analytical Jacobian matrix $\mathbf{J}(\boldsymbol{\theta}) = \frac{\partial \mathbf{p}}{\partial \boldsymbol{\theta}}$:

$$\mathbf{J}(\boldsymbol{\theta}) = \begin{bmatrix} -l_1 \sin\theta_1 - l_2 \sin(\theta_1 + \theta_2) & -l_2 \sin(\theta_1 + \theta_2) \\ l_1 \cos\theta_1 + l_2 \cos(\theta_1 + \theta_2) & l_2 \cos(\theta_1 + \theta_2) \end{bmatrix}. \tag{18}$$

To compute the local sensory-to-motor transformation under delayed feedback, we recover the estimated joint configuration $\hat{\boldsymbol{\theta}}_{i,t} = \text{IK}(\hat{\mathbf{p}}_{i,t})$ corresponding to each microzone's predicted current position $\hat{\mathbf{p}}_{i,t}$ (from Sec. 3.1.3). Using the Moore-Penrose pseudoinverse $J^\dagger = J^\top (J J^\top)^{-1}$ of the analytical Jacobian, we compute the local sensory-to-motor transformation using each microzone's predicted current position $\hat{\mathbf{p}}_{i,t}$ from Sec. 3.1.3:

$$\Delta \mathbf{u}_{i,t} = J(\text{IK}(\hat{\mathbf{p}}_{i,t}))^\dagger \, \Delta \mathbf{p}_{i,t}. \tag{19}$$

Aggregating local contributions across active microzones $\mathcal{K}_t$, the executed motor command follows:

$$\mathbf{u}_t^{\text{exec}} = \mathbf{u}_t + \sum_{i \in \mathcal{K}_t} \Delta \mathbf{u}_{i,t}. \tag{20}$$

While this pathway achieves rapid adaptation via kinematic priors, it relies on closed-form invertibility and workspace dimensionality assumptions, contrasting with the generalizable learned operator $\mathbf{M}_i$ defined in the main text.

A.6 Ablations and model variants

Table 4 summarizes the components retained in each model variant.

Table 4: **Model variants used in the ablation experiments.** Geometry denotes fixed error-direction-based routing; variants without geometry instead use a learned gate.

Variant	Geometry	Routing	Direct	Feedback
Full MoM	✓	Top-2 ($K = 2$)	✓	✓
No geometry	×	Learned	✓	✓
No direct	✓	Top-2 ($K = 2$)	×	✓
No geometry + No direct	×	Learned	×	✓
No feedback	✓	Top-2 ($K = 2$)	✓	×
Hard gate	✓	Top-1 ($K = 1$)	✓	✓
Soft gate	✓	All ($K = 4$)	✓	✓